

# On-Chip Interconnects: Circuits and Signaling from an MPSoC Perspective

Wayne Burleson

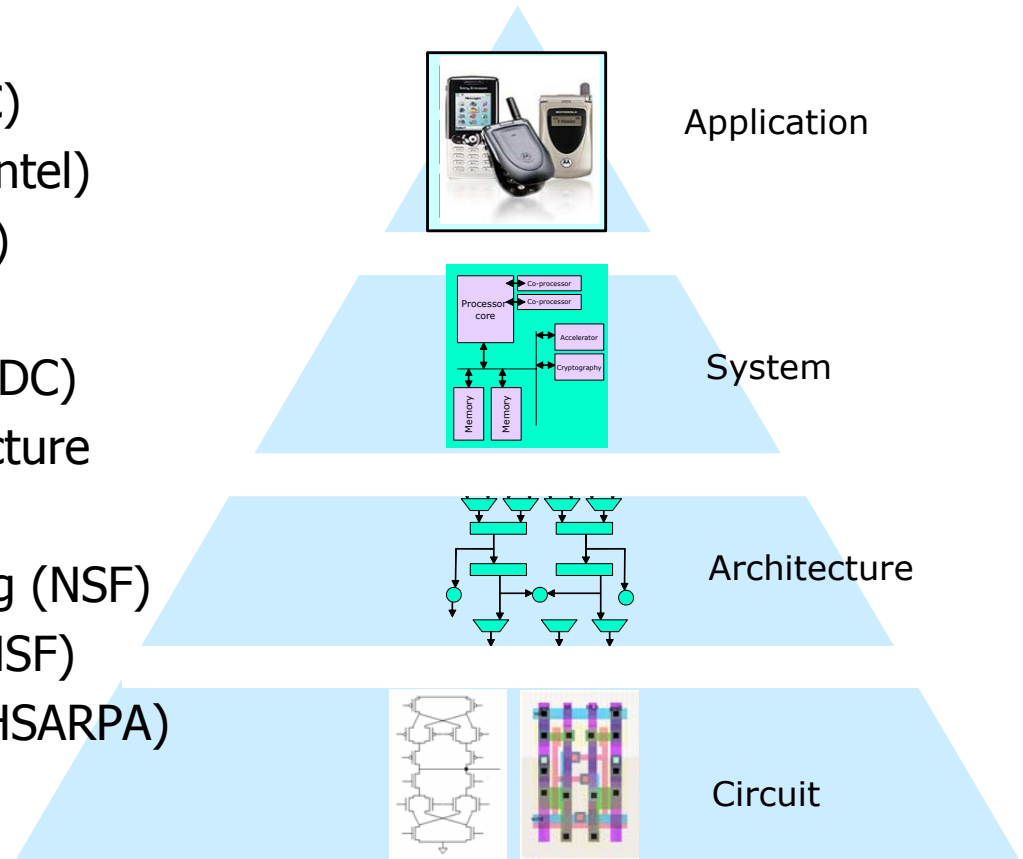
Associate Professor, Dept ECE

University of Massachusetts Amherst

[burluson@ecs.umass.edu](mailto:burluson@ecs.umass.edu)

# My Perspective

- VLSI Signal Processing, BS/MS MIT 1983, PhD Colorado 1989
- Worked as a VLSI Designer (Fairchild, VTI) and teach VLSI Design
- Research in VLSI Circuits
  - Low-power (NSF, SRC)
  - Interconnects (SRC, Intel)
  - Wave-pipelining (NSF)
  - SRAM (Intel, CRL)
  - Soft-errors (Intel, MMDC)
- Research in VLSI Architecture
  - Adaptive SOC (NSF)
  - VLSI Signal Processing (NSF)
  - Video, 3D Graphics (NSF)
  - Embedded Security (HSARPA)




# Objectives

- Provide insight into on-chip interconnects from a **circuit designer's** perspective
- Survey **recent research** in on-chip interconnects
- Present MPSoC interconnect **requirements and metrics**
- Show how to **compare** circuit and signaling solutions
- Discuss the impact of **uncertainties** (process variation, noise, temperature)
- Show **CAD** support for interconnect design and estimation
- Some examples

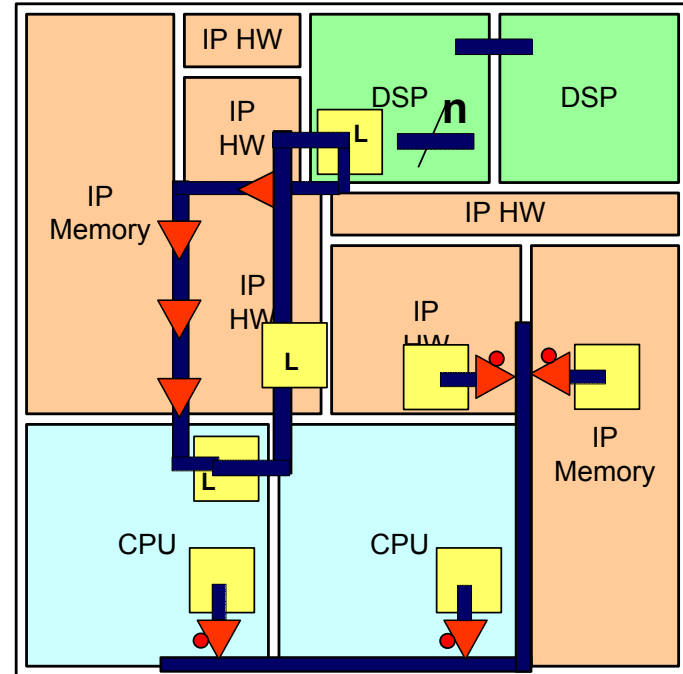
# On-Chip Interconnect: Levels of Abstraction

- Network level
  - CDMA
  - TDMA
- System level
  - Communication Links
  - Adaptive supply voltage links
- Architecture level
  - AMBA™
  - CoreConnect™
- Circuit level
  - Low Swing
  - Coding
  - Single / Differential



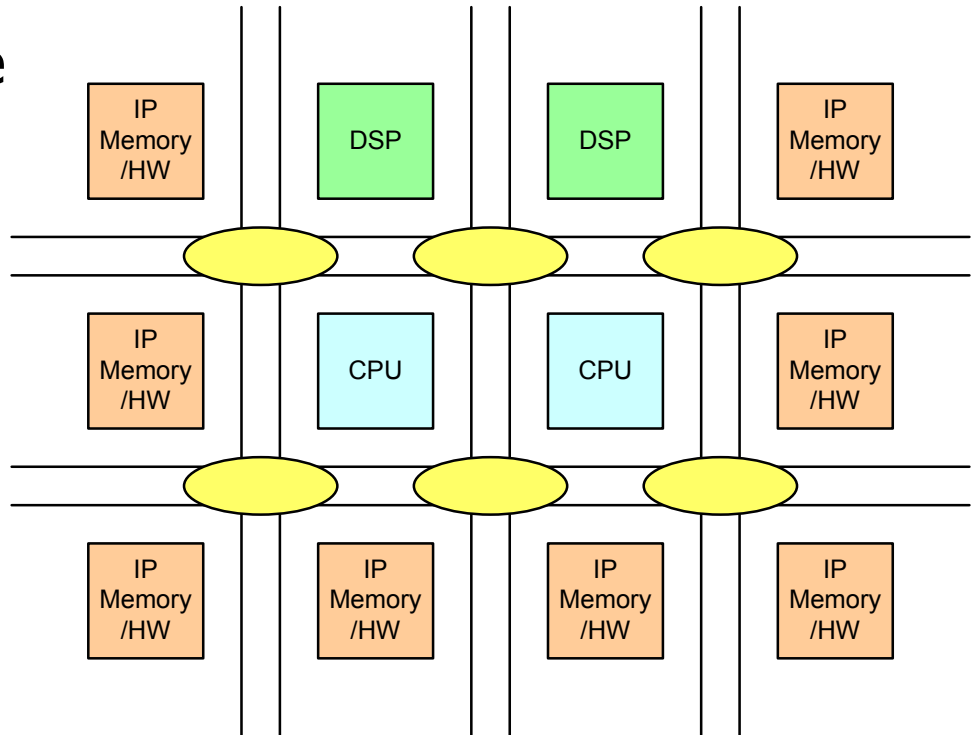
# MPSoC Interconnect Requirements

- Intra-core vs. inter-core,
- Bus-width (1,8,...32,64,...)
- Adjacent core vs. long-haul
- Repeated vs. unrepeated
- Single-cycle vs. pipelined
- Bus vs. point-to-point
- Synch vs. asynch
- Metrics:
  - Latency
  - Bandwidth
  - Noise
  - Area
  - Power/Energy



# MPSoC NoC Interconnect issues

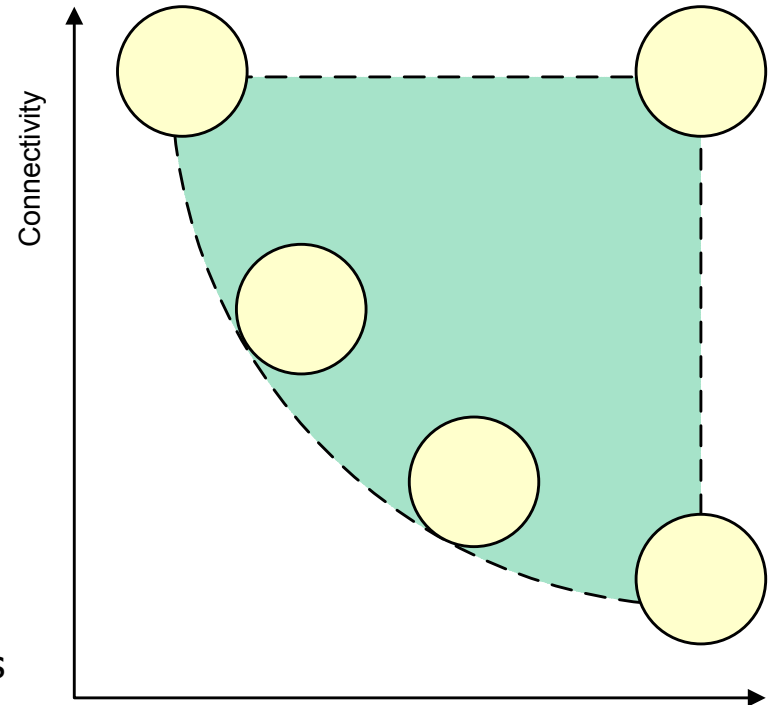
- Granularity of cores, ie wirelength
- Bus width
- Network topology
- Link layer (fault-tolerance, etc.)
- Synchronous vs. Asynchronous



Source : Ahmed Jerraya, Wayne Wolf, **Multiprocessor Systems-on-Chips**, Elsevier 2005

# MPSoC "Bus" Alternatives

- Fixed Bus [Bergamaschi, DAC, 2000]
  - Point to point communication
  - Signals between cores transferred by dedicated wires
- FPGA-like Bus [Cherepacha, FPGA Sym, 1994]
  - Programmable interconnects
  - Employ static network
- Arbitrated Bus [IDT Inc., 2000]
  - Time-shared multiple core connectivity
  - Use arbitrator
- Hierarchical Bus [AMBA, ARM Inc]
  - Combine multiple buses using bus bridges
  - Separate buses for cores and I/O
- NoC Bus [Dally, DAC, 2000]
  - Resources communicate with data packets
  - Use switch fabric



# SoC Bus Standards: CoreConnect™ and AMBA™

	IBM CoreConnect Processor Local Bus	ARM AMBA 2.0 AMBA High-performance Bus
<b>Bus Architecture</b>	32-, 64-, and 128-bits Extendable to 256-bits	32-, 64-, and 128-bits
<b>Data Buses</b>	Separate Read and Write	Separate Read and Write
<b>Key Capabilities</b>	Multiple Bus Masters 4 Deep Read Pipelining 2 Deep Write Pipelining Split Transactions Burst Transfers Line Transfers	Multiple Bus Masters Pipelining Split Transactions Burst Transfers Line Transfers
	On-Chip Peripheral Bus	AMBA Advanced Peripheral Bus
<b>Masters Supported</b>	Supports Multiple Masters	Single Master: The APB Bridge
<b>Bridge Function</b>	Master on PLB or OPB	APB Master Only
<b>Data Buses</b>	Separate Read and Write	Separate or 3-state



# Interconnect Geometry Scaling

Layer	Pitch (nm)	Thick (nm)	AspectRatio
Isolation	220	320	-
Polysilicon	220	90	-
Contacted gate pitch	220	-	-
Metal 1	210	170	1.6
Metal 2	210	190	1.8
Metal 3	220	200	1.8
Metal 4	280	250	1.8
Metal 5	330	300	1.8
Metal 6	480	430	1.8
Metal 7	720	650	1.8
Metal 8	1080	975	1.8


**65nm Intel<sup>®</sup> Technology**

LAYER	PITCH	THICK	AR
Isolation	240	400	-
Poly-Si	260	140	-
Metal 1	220	150	1.4
Metal 2,3	320	256	1.6
Metal 4	400	320	1.6
Metal 5	480	384	1.6
Metal 6	720	576	1.6
Metal 7	1080	972	1.8

**90nm Intel<sup>®</sup> Technology**

- Weak scaling of vertical dimension compared to horizontal dimension
- Extremely high height/width aspect ratios
- Reduces degradation of interconnect resistance

# Interconnect RC Trend

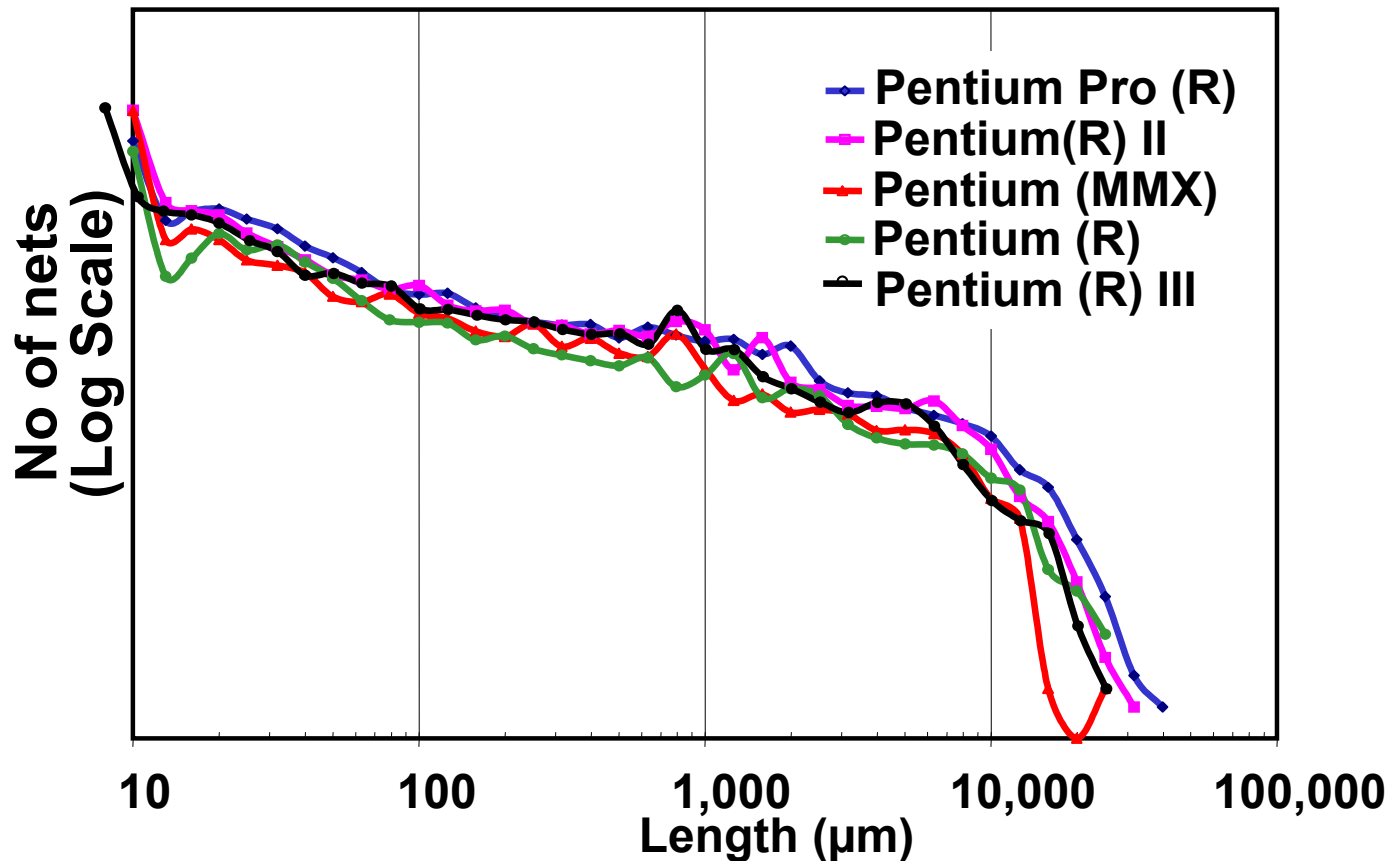


% increase each generation			
	R	C	RC
Poly	45%	-2%	42%
M1	53%	5%	61%
M2	46%	12%	62%
M3	39%	8%	51%
M4	18%	24%	46%

λ **RC/μm increases 40-60% per generation**

λ **Copper, low-K dielectric: modest benefit**

# Interconnect Distribution Trend

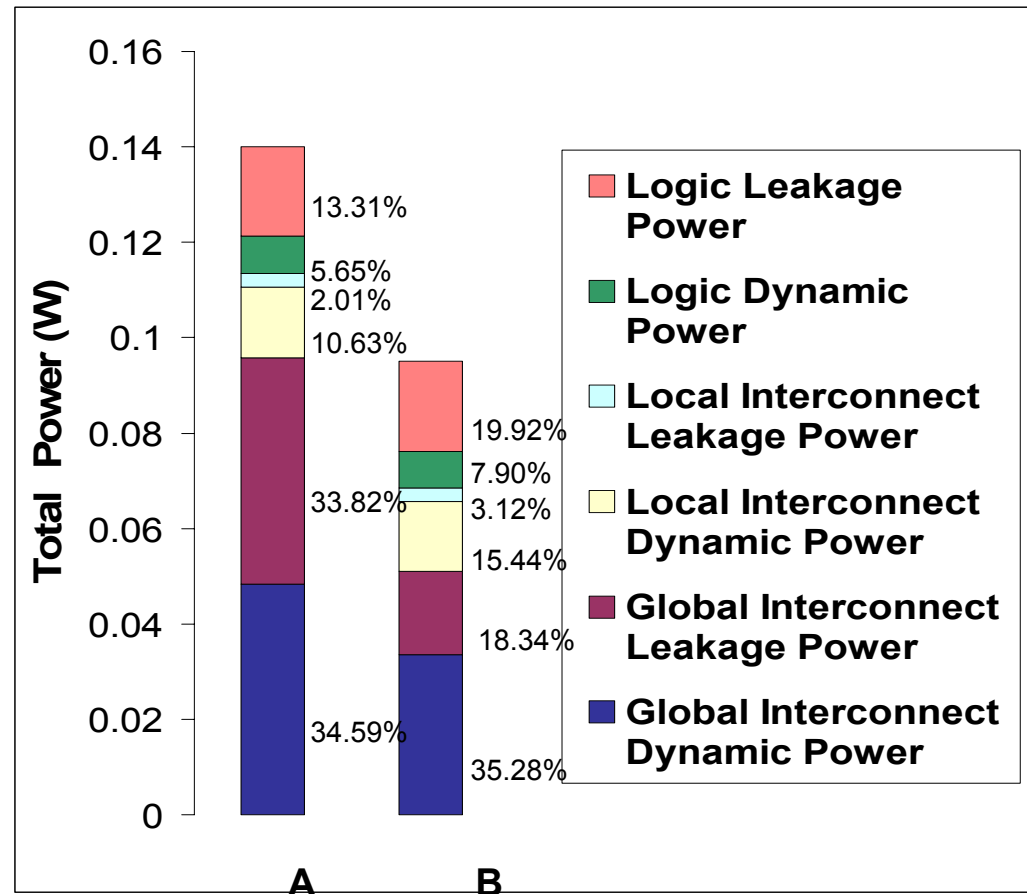


**RC/ $\mu\text{m}$  scaling trend is only one side of the story...**  
**Average wire lengths don't scale well**


# Interconnect Power Consumption

- Using Vdd programmability
- High Vdd to devices on critical path
- Low Vdd to devices on non-critical paths
- Vdd Off for inactive paths

A – Baseline Fabric  
 B – Fabric with Vdd Configurable Interconnect



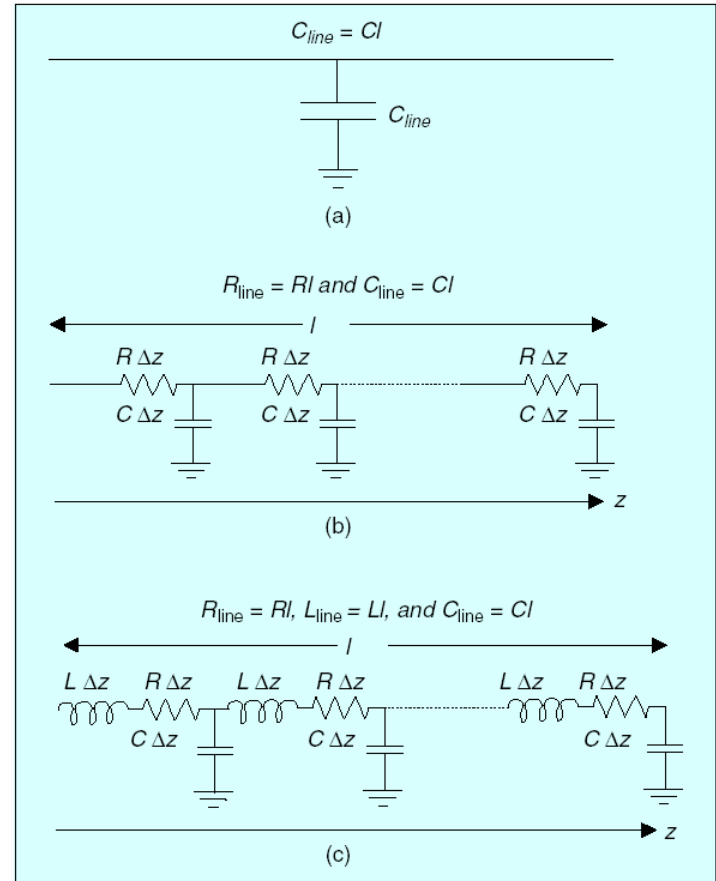
# Interconnect Modeling




a. Capacitive model

b. RC model

c. RLC model



# Interconnect Issues – Signal Integrity

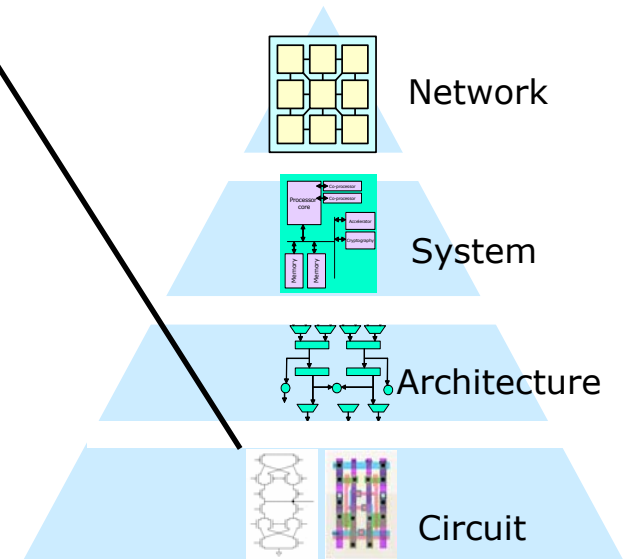


Ismail, TVLSI, 2002


- Inductance becoming important
- Self inductance results in ringing
- Mutual inductance results in crosstalk

# Circuit and Signaling Solutions

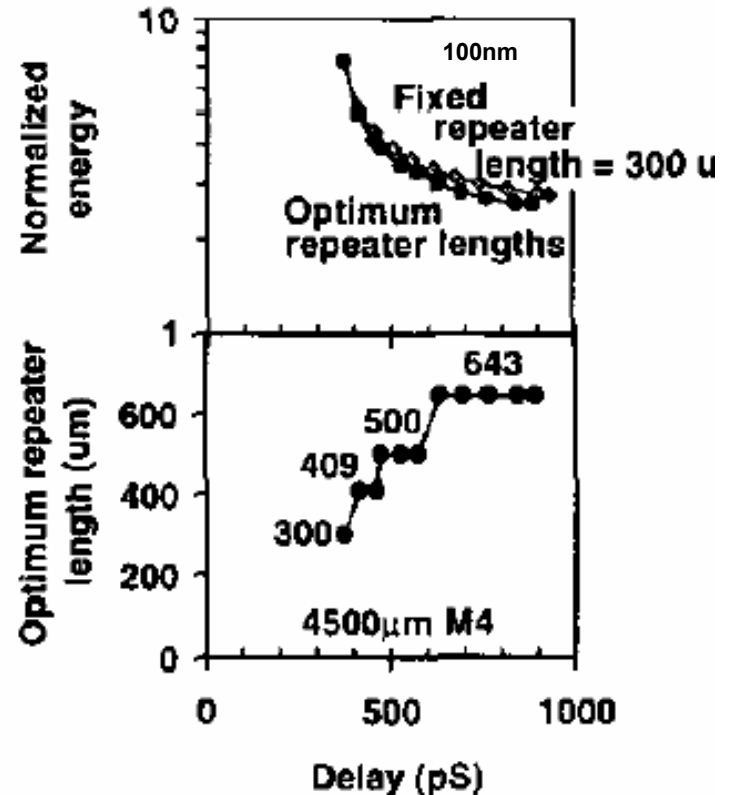
- Conventional Circuit techniques
  - Repeater insertion
  - Booster insertion
- Low Swing techniques
  - Pseudo differential interconnect
  - Differential Current sensing
- Bus encoding techniques
  - Transition aware encoding
  - Low Power encoding for crosstalk reduction
- Signaling techniques
  - Multi-level signaling
  - Near speed of light signaling



# Interconnect Circuits - Repeater




- Optimum repeater insertion reduces interconnect delay.
- Optimized energy-delay tradeoffs used to satisfy design criteria.






# Interconnect Circuits - Repeaters

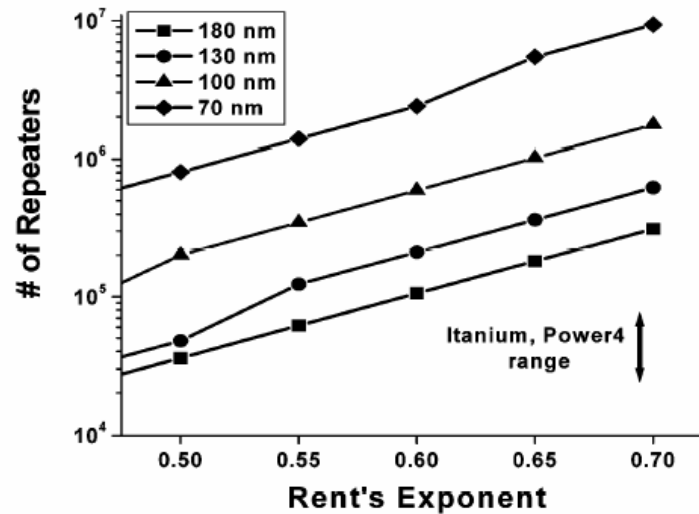


Maheshwari, PhD Thesis, 2004

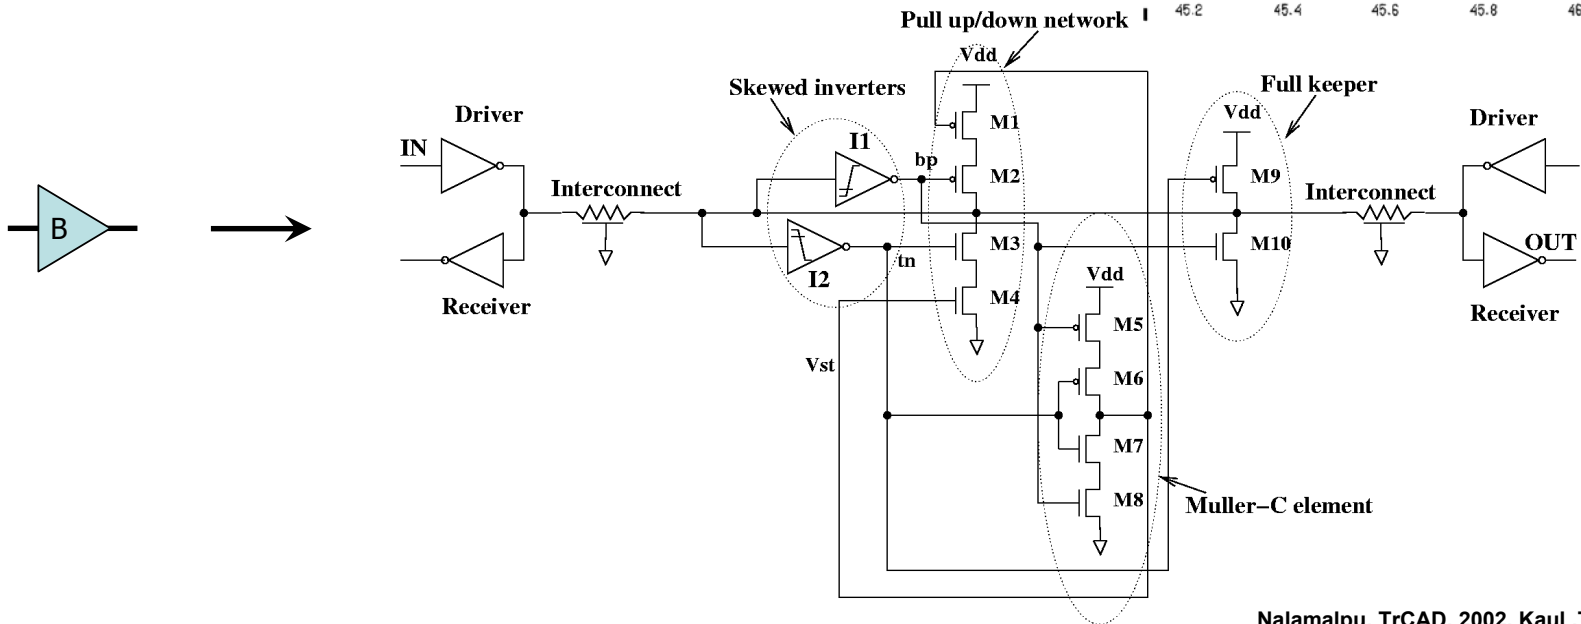




Dynamic Power in Repeated Wires

Magen, SLIP, 2004




# Interconnect Circuits – Boosters



# Interconnect Circuits - Low swing

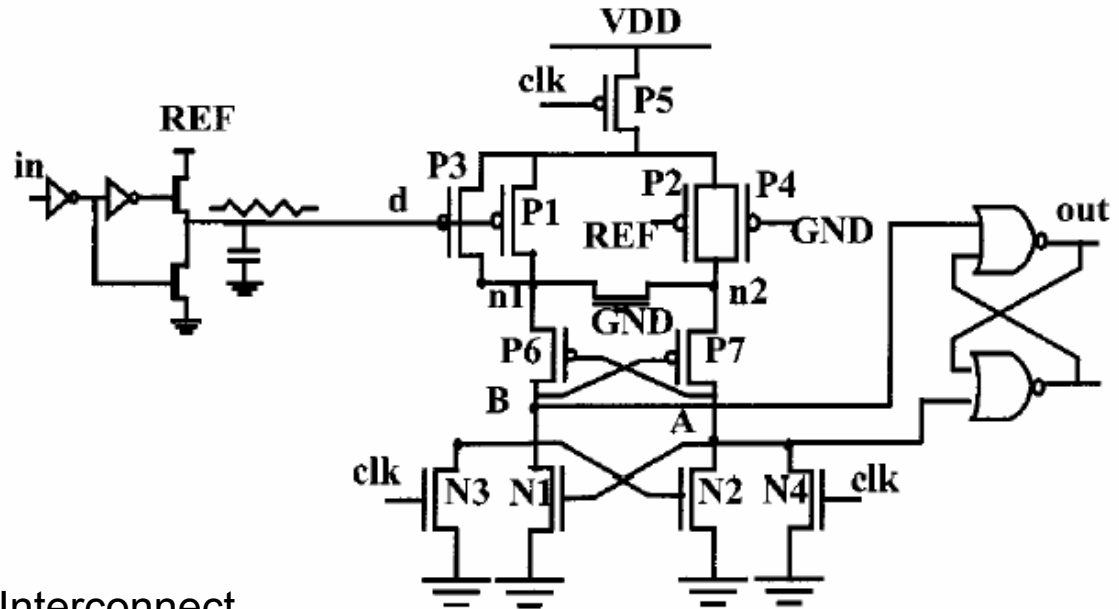
- Reduce the swing on the interconnect
- Use a PMOS/NMOS device to provide a resistive path
- Reduces dynamic power since interconnect is not charged full rail.
- Noise immunity low at the output due to reduced swing
- Static power dissipation due to low impedance path



- Current mode, voltage mode
- Single ended, differential

# Interconnect Circuits – Low Swing

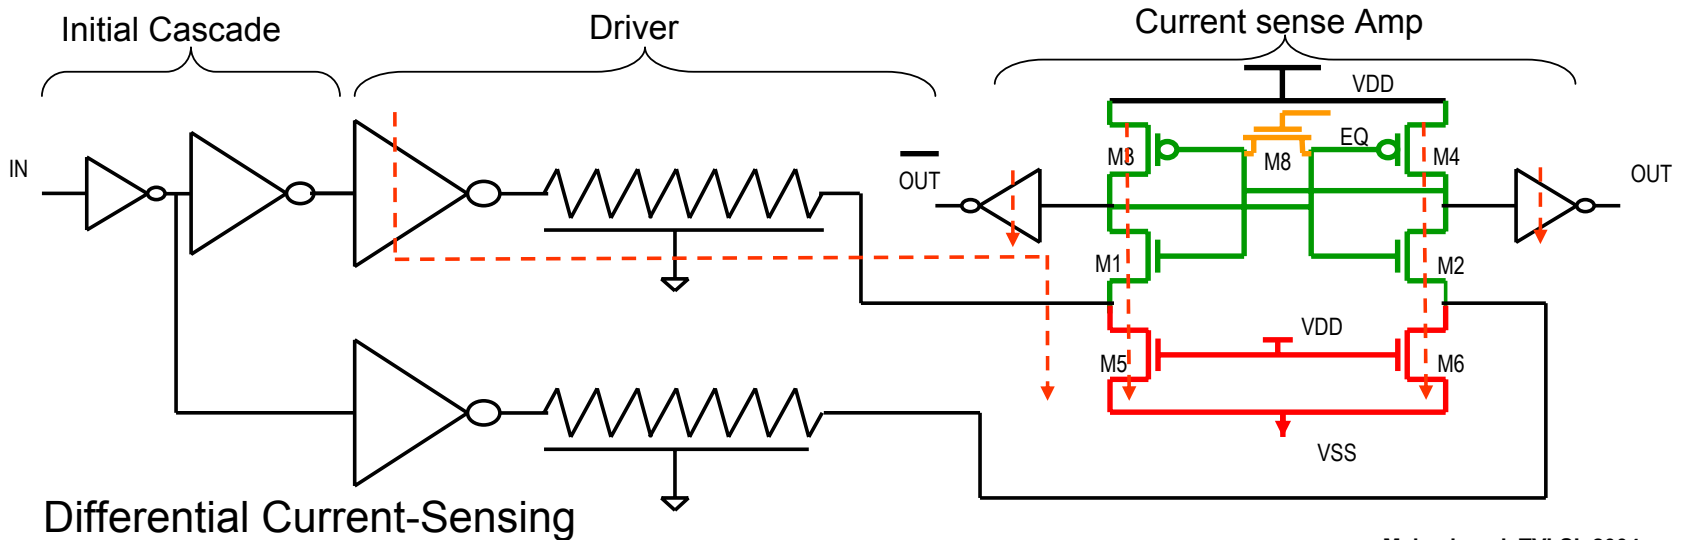
- Voltage mode
- Single ended
- One wire per bit
- Receiver not sensitive to supply variation



Pseudo Differential Interconnect

Schemes	Energy (PJ)			Delay (ns)			E•D (PJ•ns)	Swing (V)	Complexity
	Driver/Wire	Receiver	Total	Driver/Wire	Receiver	Total			
CMOS	11.45	0.15	11.6	1.64	0.47	2.11	24.5	2.0	least area overhead
PDIF	1.32	0.60	1.92	1.65	0.75	2.40	4.6	0.5	timing, 1 REF

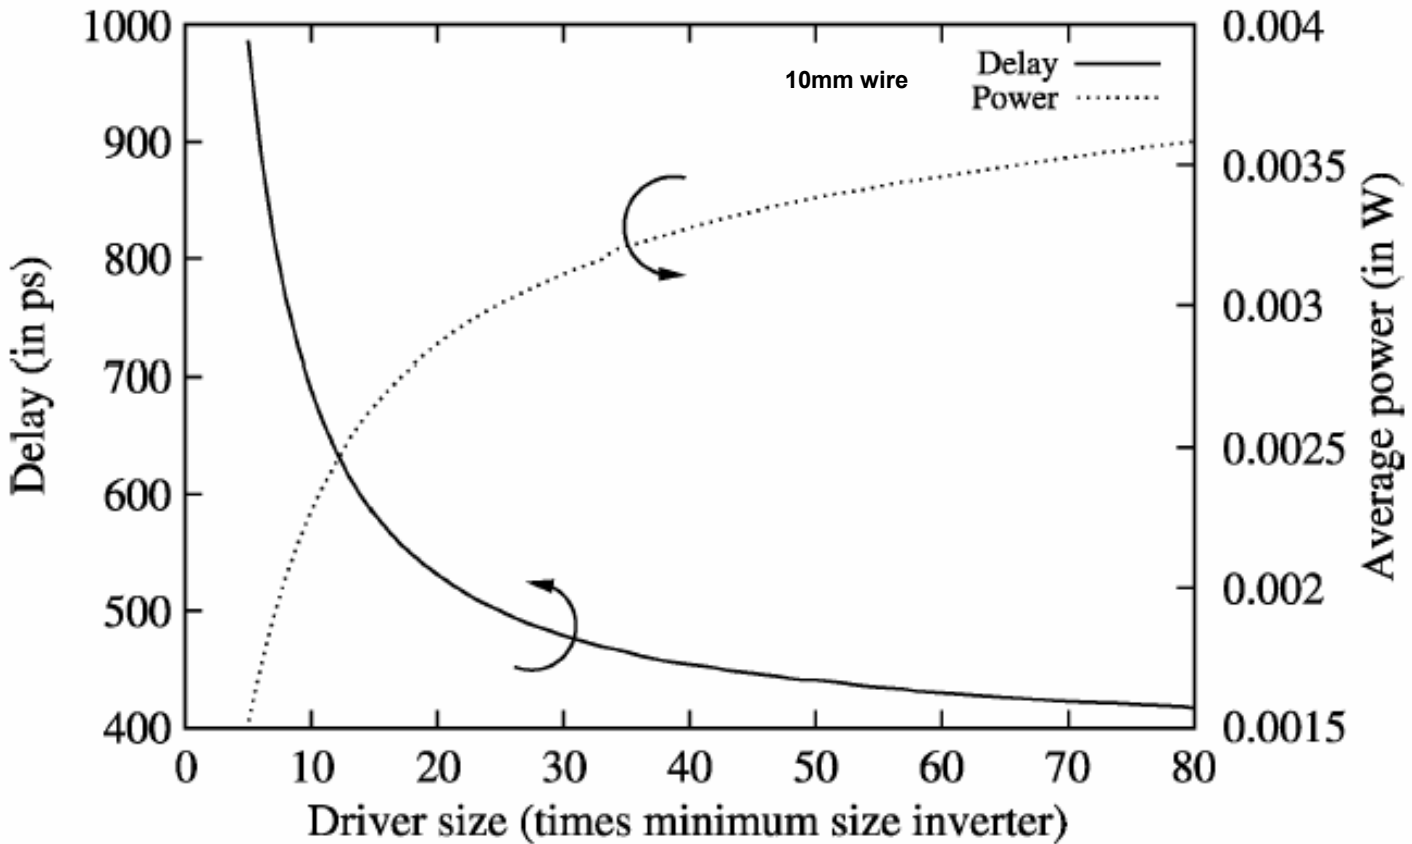
# Interconnect Circuits – Low Swing



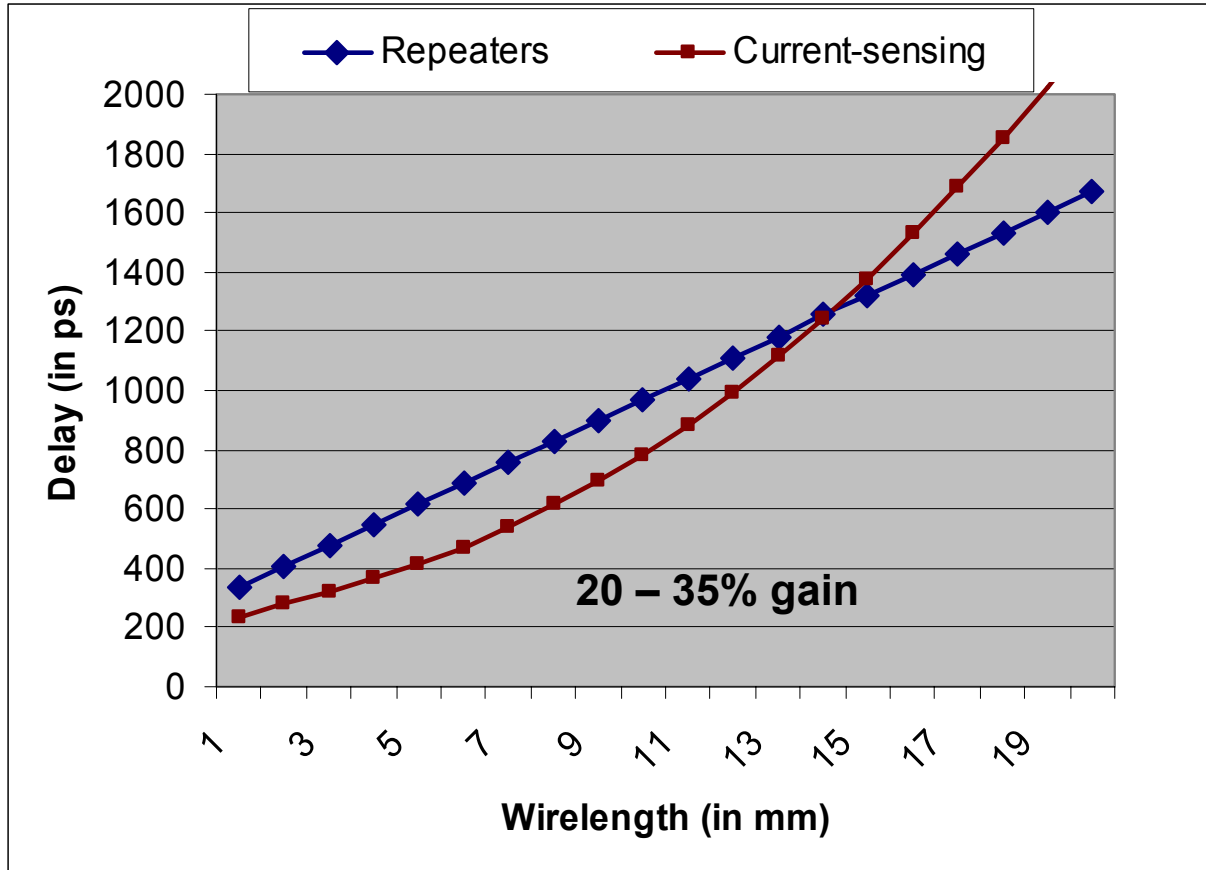
Maheshwari, TVLSI, 2004

- Current mode, Differential
- Avoids charging and discharging wire capacitance
- No repeaters along the wire: Avoids placement constraints
- Suffers from static power dissipation (paths shown by dashed lines)

# Delay-Power tradeoffs

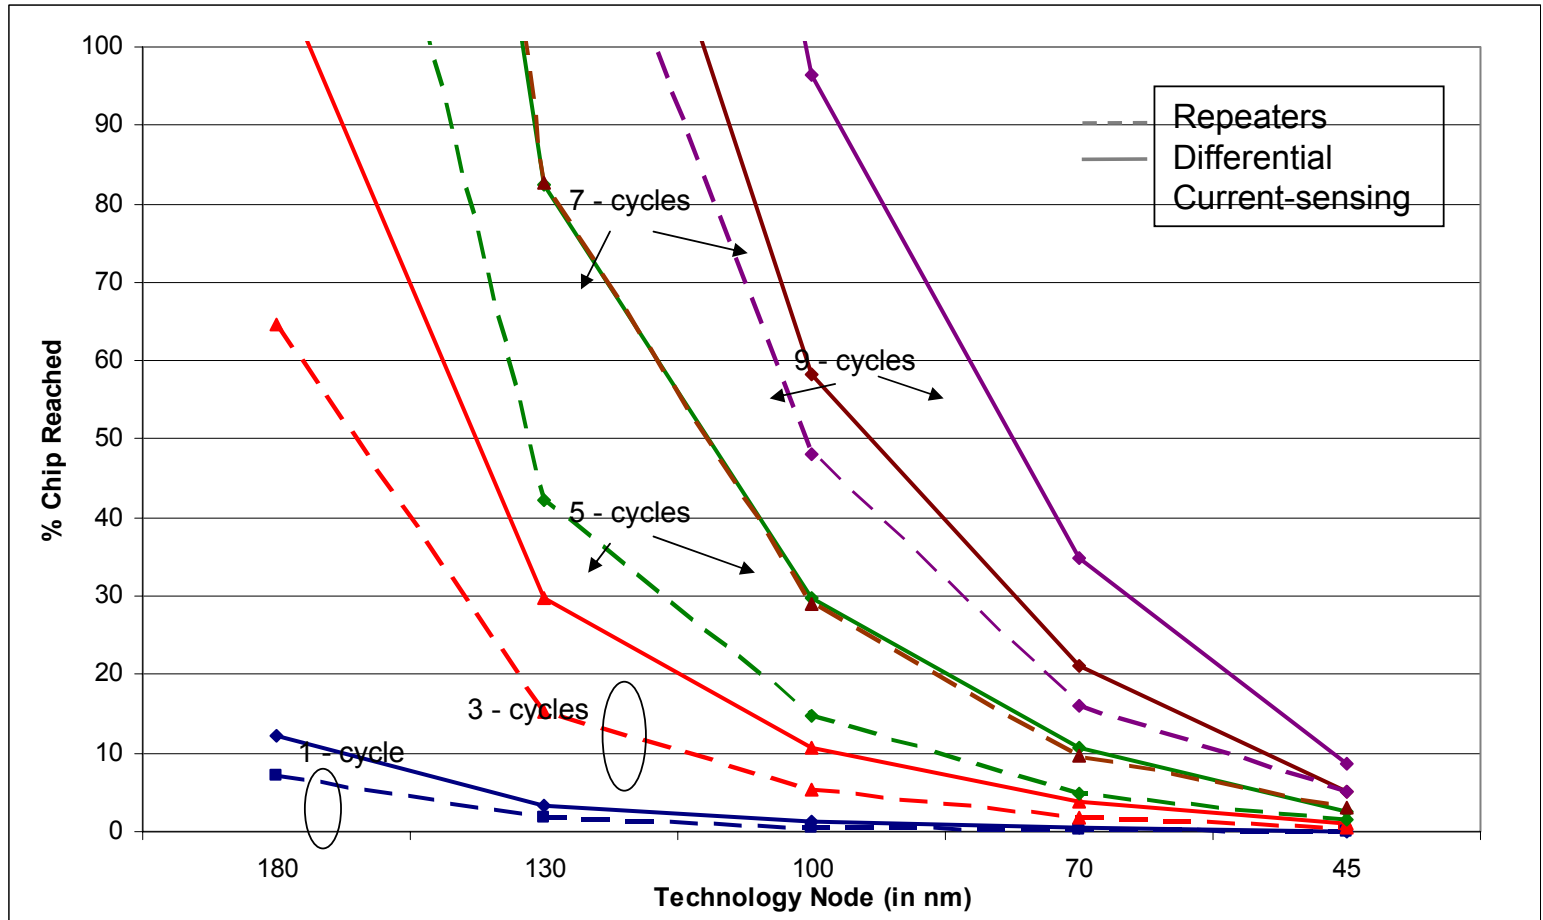


# Delay vs. wirelength



Intel 90nm (wires with 2x min. width)

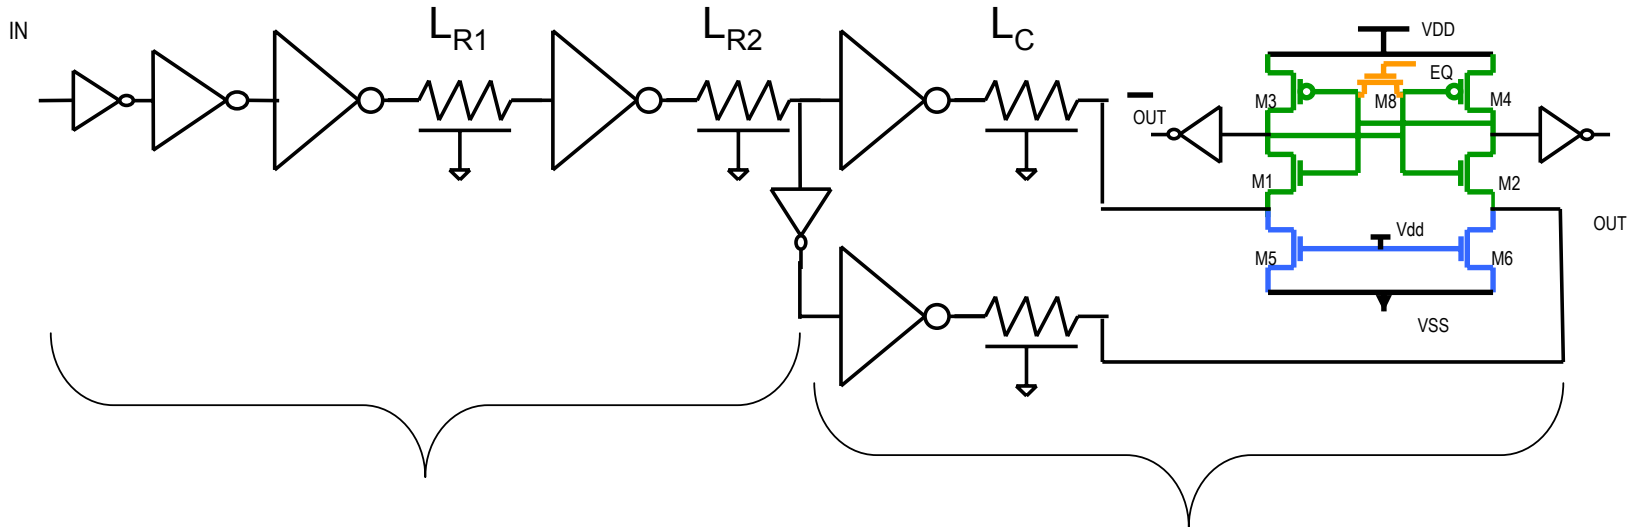
# % chip coverage in n cycles



Percentage of Chip Coverage



# Hybrid Repeaters & Current-sensing



Uniform Repeater Insertion

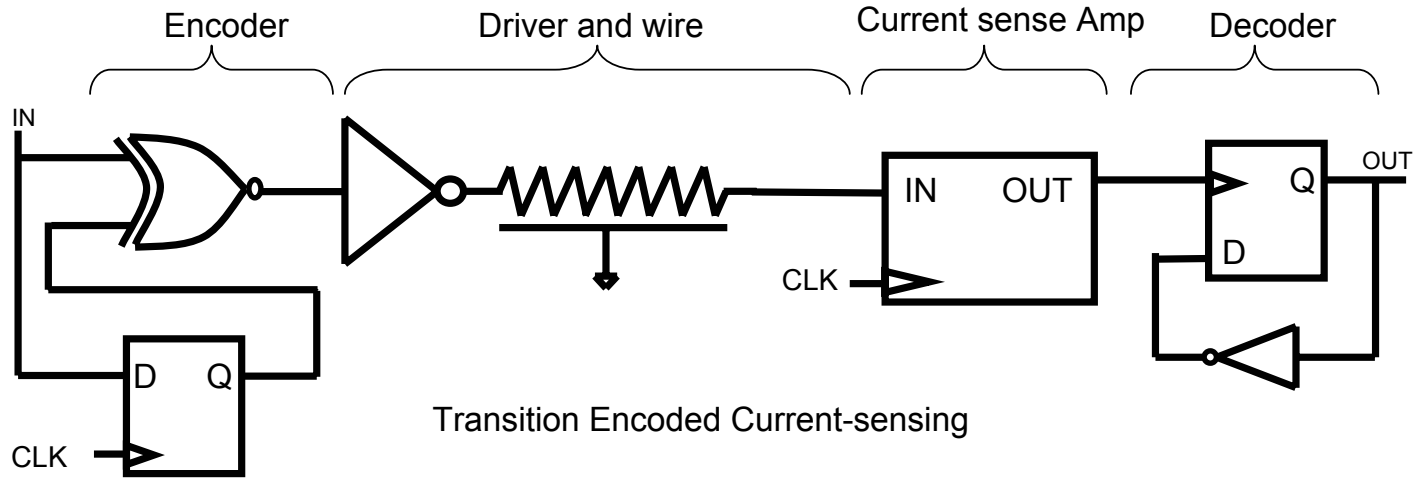
Current-sensing e.g. DCS

$$\sum_{i=1}^n L_{Ri} + L_C = L$$


$$\frac{\sum_{i=1}^n L_{Ri}}{L} = ?$$

How much wire driven by repeaters ?

# Eliminating bus static power dissipation



Transition Encoded Current-sensing




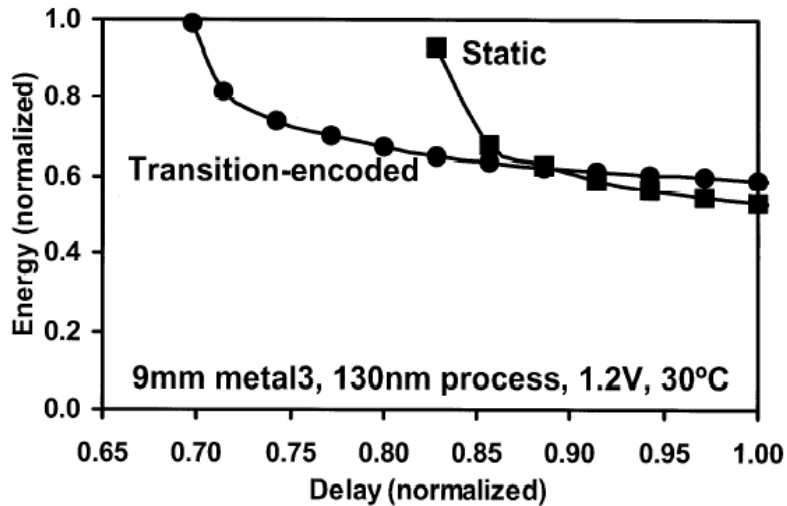
Current-pulse signaling

- Send current only when there is a transition
- Hold the bus at GND otherwise
- Encoder and decoder overhead


# Interconnect Solutions - Bus Encoding

- Reduce dynamic power due to switching activity on a bus
  - Transition encoding, spatial encoding, invert encoding, pattern encoding
- Various encoding target different aspect of interconnect
  - Delay, power, energy, crosstalk, area
- Cost of encoding/decoding
  - Power, area, latency, additional wires

# Interconnect Solutions – Bus encoding



- Uses a dynamic bus configuration
- Encoder translates input transition activity into an output logic state
- Decoder uses encoded signal to reconstruct the original input using its stored state information to distinguish between the two input transitions.

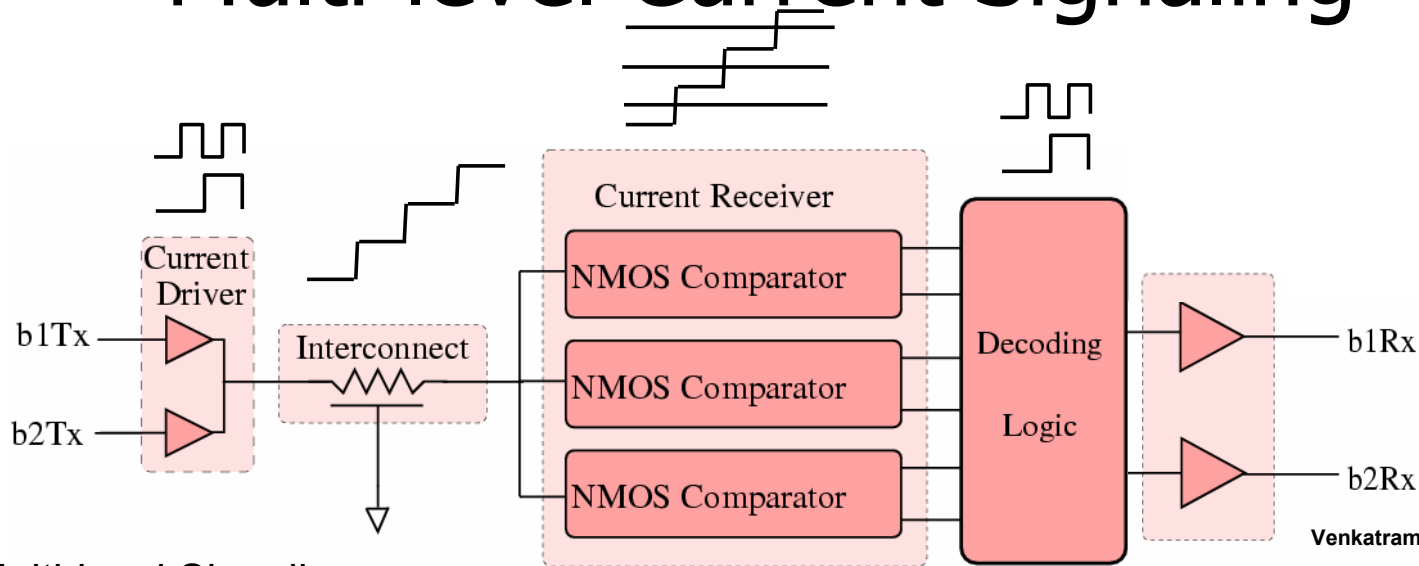


Transition Encoded Dynamic Bus

# Interconnect Solutions - Bus Encoding


- Bus invert encoding
  - Checks each cycle if there is a possibility of greater than 50% transitions on the bus
  - Decides whether sending the true or compliment form of the signals
  - Reduces the switching activity
  - Requires one additional wire to inform receiver whether the bus is true or complement
  - Numerous extensions and improvements for different statistical assumptions and metrics

# Multi-level Current Signaling

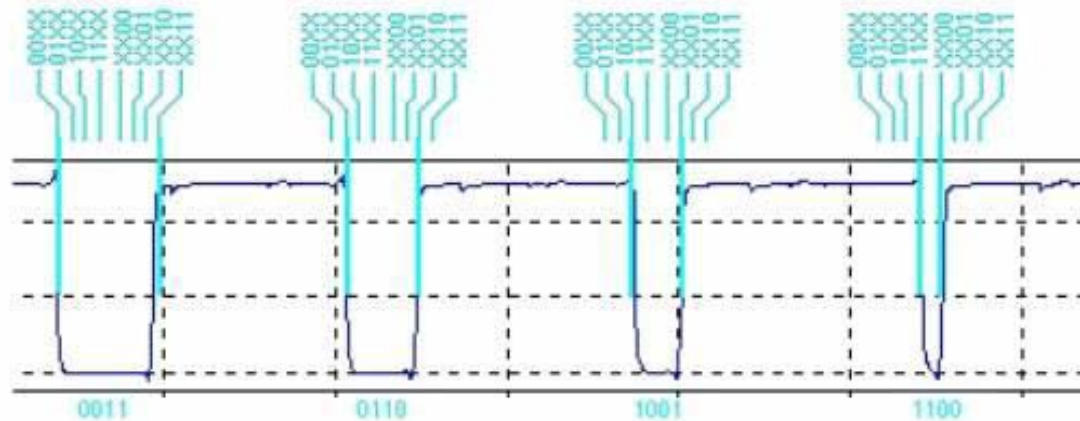


## Multi-level Signaling

- Encode two or more data bits and transmit on interconnect.
- The two or more data bits are encoded into four or more current levels. Current provides more head-room than voltage!
- Sense the current levels and decode the original signals

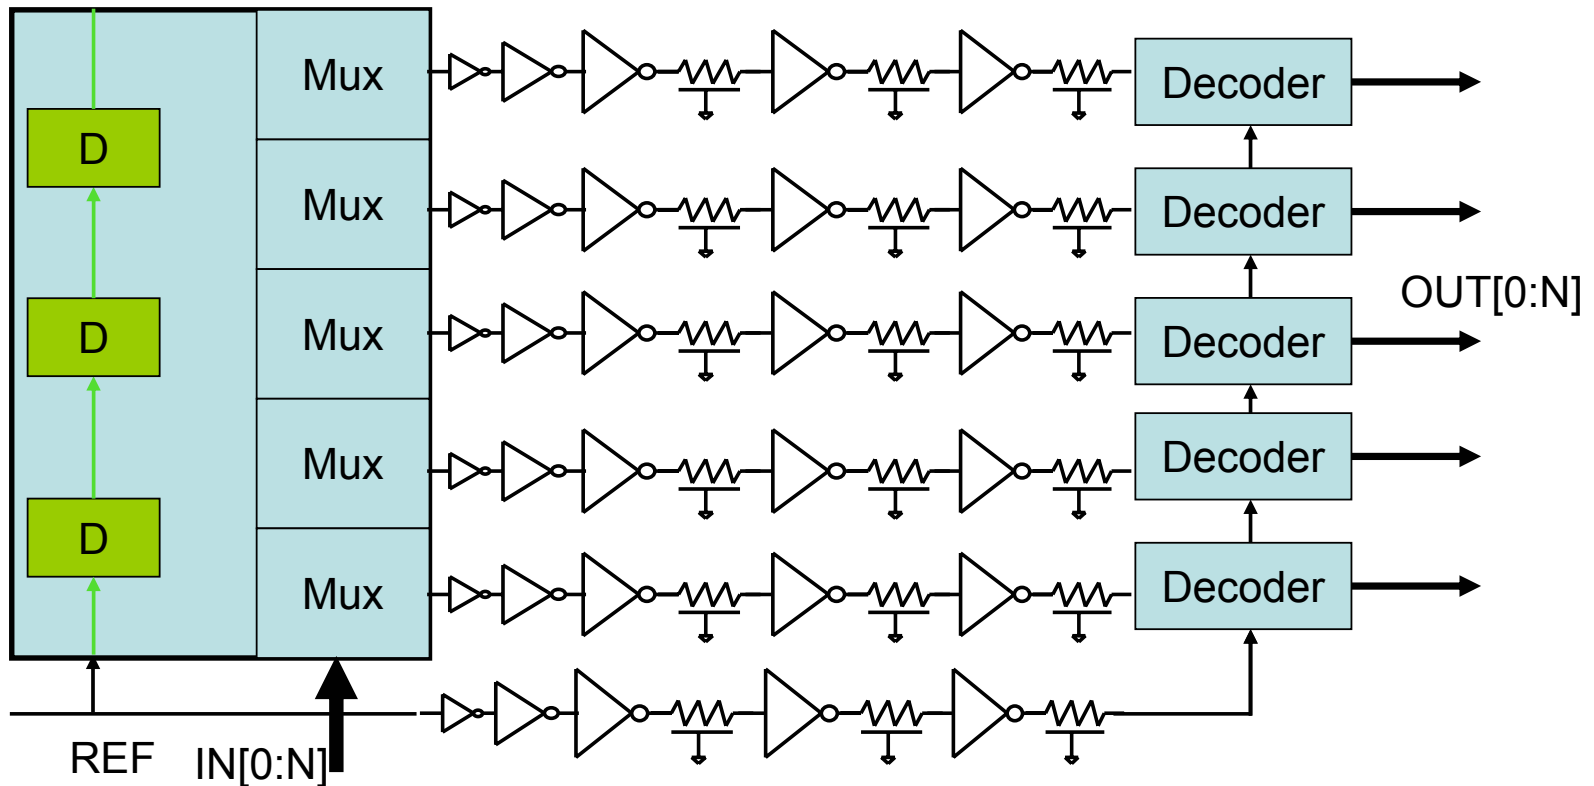


# Phase Coding



- Actually phase modulation
- Transmitting multiple bits in one transition
  - Significant power and area savings
  - Increased bandwidth
- Phase coding – Phase determines the data
- How to deal with timing uncertainty?

# Open Loop Phase Coding



- Delay elements can be shared across wires
- Supply noise, Process variation etc. can result in errors



# Measured Results: Closed Loop

- 16-bit 5mm long bus, 0.27u wide, 0.27u spacing, shielded, 1GHz
- Repeater insertion, Transition encoding used
- Encode in  $\frac{1}{2}$  cycle and use  $\frac{1}{2}$  cycle for decode

Encoding Levels (bits/wire)	Encoder Overhead (mW)	Decoder Overhead (mW)	Phase coding power (mW)	Repeater bus (mW)
2	0.33	1.00	5.61	8.56
3	0.47	1.33	5.01	8.56
4	0.62	1.52	4.28	8.56 <sub>33</sub>

# Near Speed of light Signaling

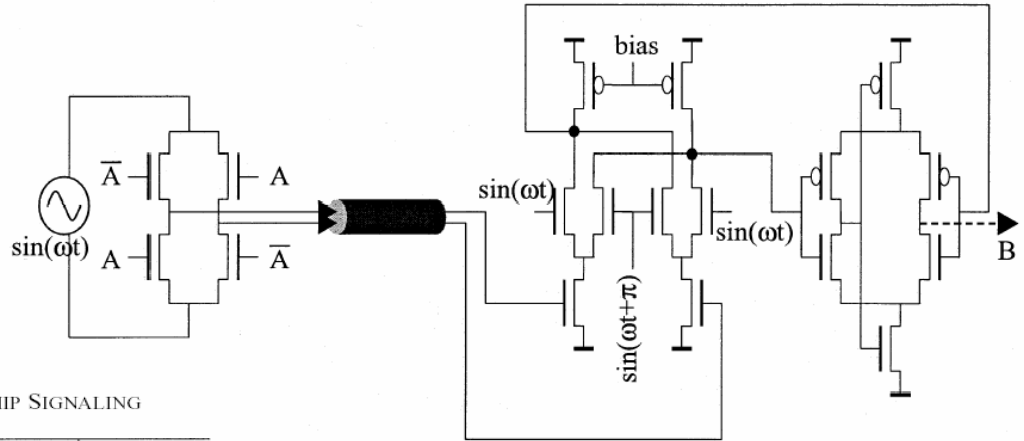



TABLE II  
PERFORMANCE OF DIFFERENT APPROACHES TO ON-CHIP SIGNALING

Signaling	Propagation Medium	Time of Flight for 20mm [ps]	Delay for 20 mm [ps]	Power [mW]	Comments
Modulation (This work)	Wide metal wires ( $\epsilon_r = 4$ )	133	300	16	Large metal area
Repeaters	Min. sized metal wires ( $\epsilon_r = 4$ )	133	1400	30	Slow
	Wide metal wires ( $\epsilon_r = 4$ )	133	400	50	Large metal area, High power
Optics (Edge-Emitting) [11]	Air ( $\epsilon_r = 4$ )	66	300	80	Packaging, Integration issues
	On-Chip Waveguide ( $\epsilon_r = 11.7$ )	228	500	80	Integration issues
Optics (VCSEL) [11]	Air ( $\epsilon_r = 1$ )	66	400	60	Packaging, Integration issues
	On-Chip Waveguide ( $\epsilon_r = 11.7$ )	228	600	60	Integration issues



- 283ps for 20mm 16um wide AL wire in 0.18um CMOS tech
- Very wide,  $R \sim 0$
- Uses frequency modulation

# Uncertainty - Process Variations



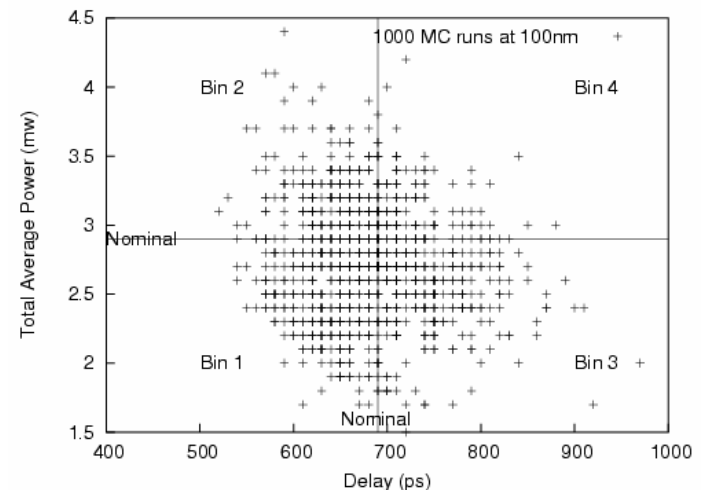
Nassif, ISQED, 2000

# Impact of Uncertainty on Delay and Power





- 100nm technology
- 1000 Monte Carlo Runs
- Power variability of 43.64%
- Delay variability of 28.95%

• 100nm Technology	
• Bin 1(High Performance) Yield	– 36.1%
• Bin 2(Low Delay) Yield	– 27.3%
• Bin 3(Low Power) Yield	– 25.1%
• Bin 4(Low Performance) Yield	– 11.5%





# Uncertainty - Temperature Variations



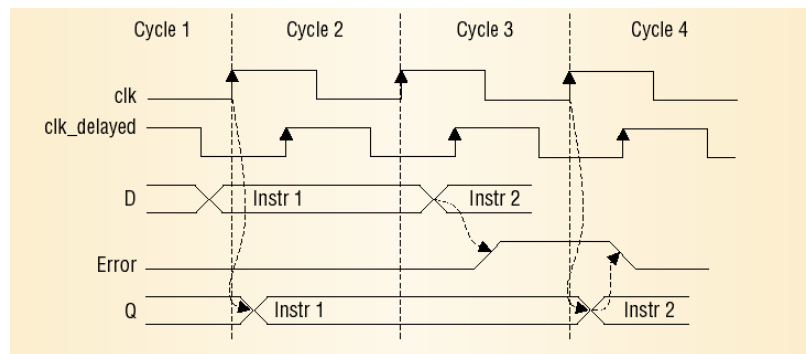
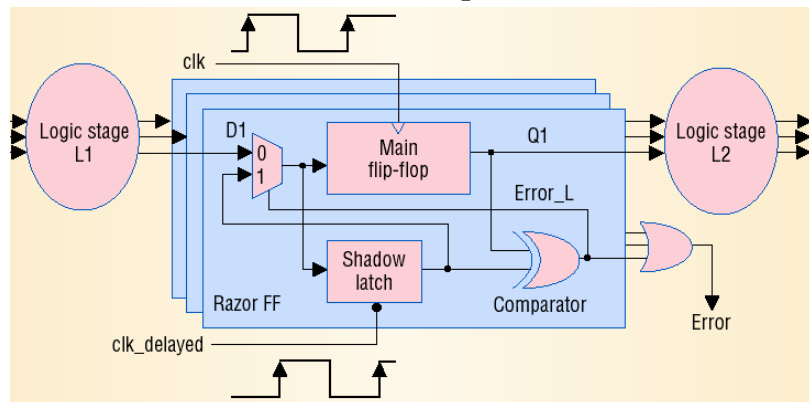
	70nm	45nm
$\frac{P_{leakage,130^\circ C}}{P_{total,130^\circ C}}$	1.5%	2.6%
$\frac{P_{dyn,130^\circ C}}{P_{dyn,30^\circ C}}$	13 %	38%
$\frac{P_{leak,130^\circ C}}{P_{leak,30^\circ C}}$	> 17	> 27

- Leakage power significantly increases with temp. in 45 nm node.
  - ◇ parabolic curvature ( $y = Cx^2$ ) in terms of varying temperatures.
- Higher temperature sensitivity (45 nm) on delay, power and leakage.
  - ◇ Accurate RLC modelization to provide underestimation.

# Uncertainty due to Power Supply Noise




# Uncertainty – Variation-tolerant Design



## • Razor methodology


- A voltage-scaling methodology based on real-time detection and correction of circuit timing errors
- Allows for energy tuning of microprocessor pipeline
- Application of Razor methodology results in up to 64% energy savings with less than 3% delay penalty for error recovery

# The Delay, Energy, Error space





# Interconnect test chips

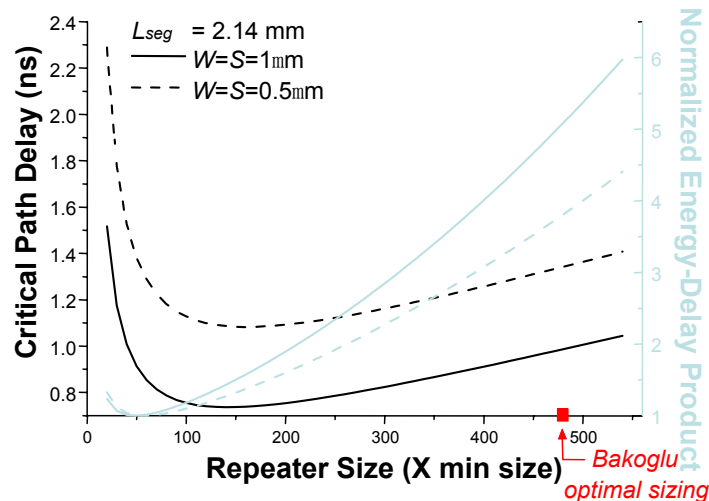


# CAD Support

- GTX (SRC-MARCO)
  - A GSRC Technology Extrapolation System
- NoCIC (UMASS)
  - SPICE-based Interconnect Calculator for aggressive Circuit techniques (current-sensing, multi-bit sensing, boosters, etc.)

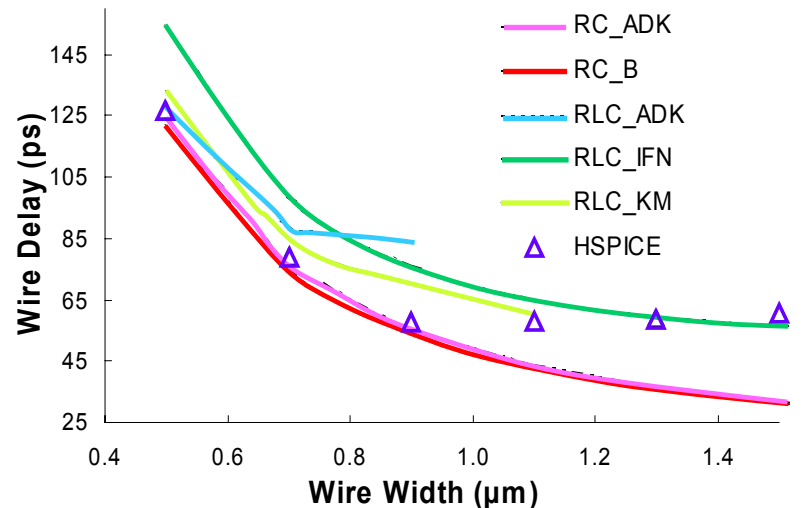

# Repeater Optimization using GTX

- Most commonly cited optimal buffer sizing expression (Bakoglu)
- In GTX:
  - Sweep repeater size for single stage in the chain
  - Examine both delay and energy-delay product



# Inductance analysis using GTX

- Five different models implemented in GTX
  - Bakoglu's model (RC\_B)
  - [Alpert, Devgan and Kashyap, ISPD 2000] (RC\_ADK)
  - [Ismail, Friedman and Neves, TCAD 19(1), 2000] (RLC\_IFN)
  - [Kahng and Muddu, TCAD 1997] (RLC\_KM)
  - Extension of [Alpert, Devgan and Kashyap, ISPD 2000] (RLC\_ADK)



# A snapshot of NoCIC

New Page 1 - Microsoft Internet Explorer provided by Comcast  
 Address: C:\Documents and Settings\Ishak Venkatraman\Desktop\Slip04\_presentation\FINAL\_NoC.htm

## NoCIC : Network-on-Chip Interconnect Calculator

### PARAMETER SELECTION WINDOW

**NoC Parameters**

Tile Size: 4mm

Bus Size: 8-bit

Supply Voltage: 1.8v

---

**Interconnect Parameters**

Signaling Technique: Repeater

Technology: 180nm

Shielding: Unshielded

---

**Analysis**

Output: Delay

Compare:

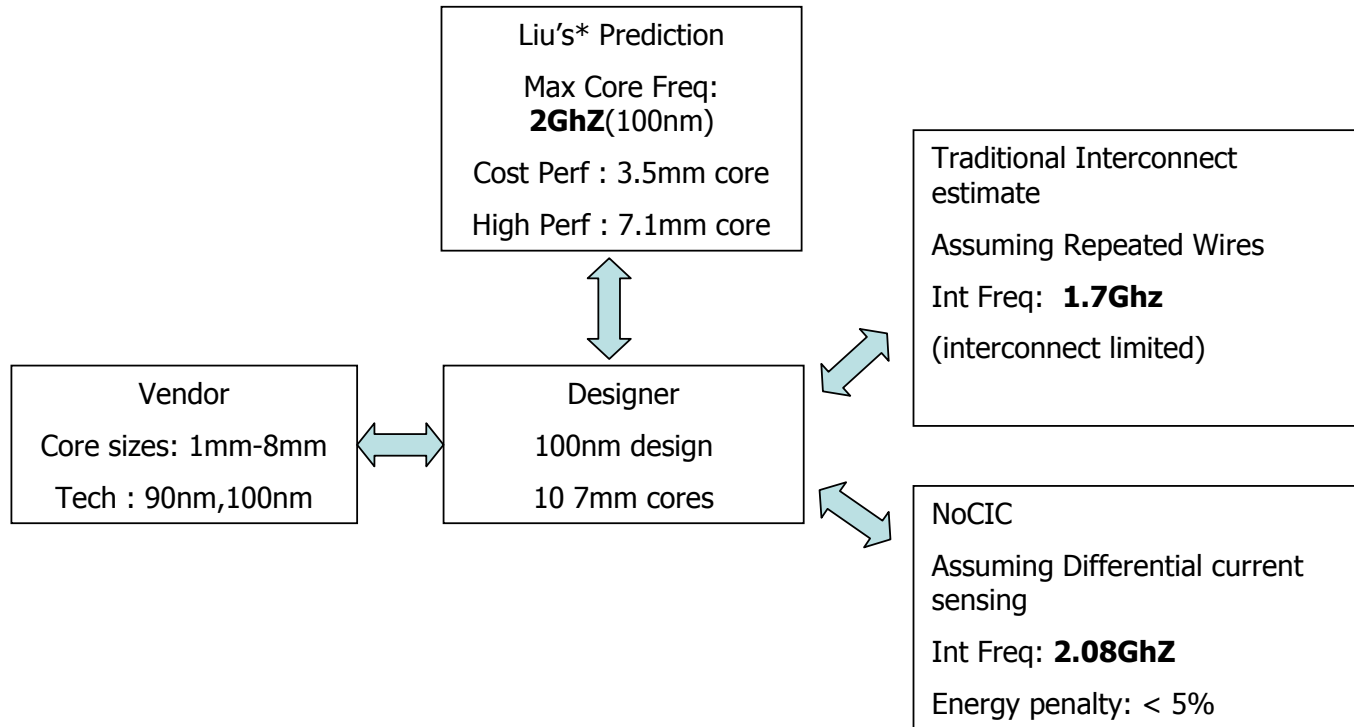
- Repeater
- Booster
- Differential Current Sensing
- Multi-level Current Signaling

Submit Reset

### PLOT DISPLAY WINDOW

Tile Size (in mm)	Repeater 180nm (ns)	DCS 180nm (ns)	Booster 180nm (ns)	NLOS 180nm (ns)
0	0	0	0	0
2	~50	~100	~80	~60
4	~100	~200	~150	~120
6	~150	~350	~250	~200
8	~200	~500	~350	~300
10	~250	~700	~450	~400
12	~300	~900	~550	~500
14	~350	~1100	~650	~600
16	~400	~1300	~750	~700
18	~450	~1500	~850	~800
20	~500	~1700	~950	~900

# Case Study : MPSoC with NOC

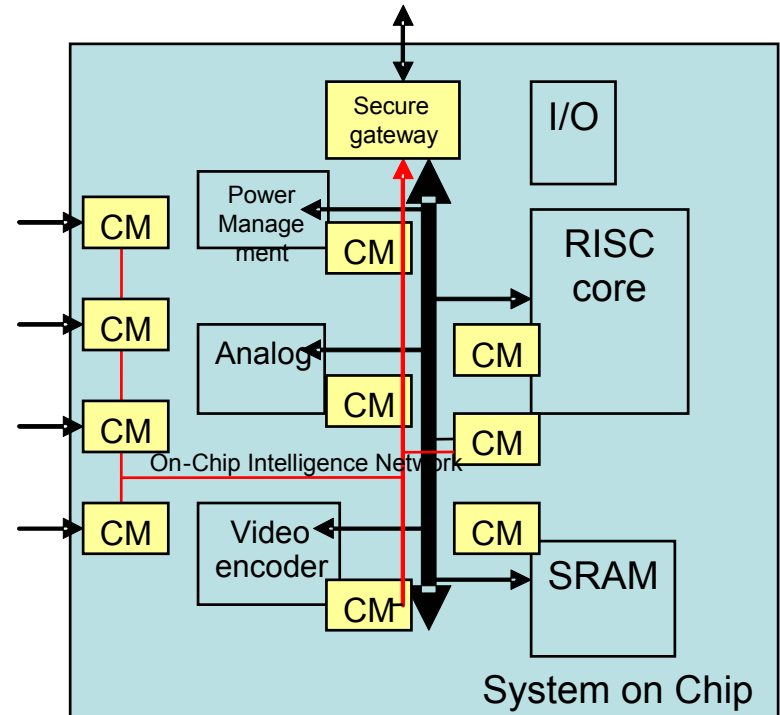


\* J. Liu et.al System level interconnect design for network-on-chip interconnect IPs, *in proceedings of the international workshop on System level interconnect prediction, SLIP 2003.*

# Case-Study: On-Chip Security

(Burleson, Tessier, Gong, Wolf, Gogniat, 2005)

- On-chip monitoring and security bus
- Latency-critical for fast detection and mitigation of attacks
- Improved power, performance and security over software-based defenses



CM = Configurable Monitor  
 OCIN = On-Chip Intelligence Network

# Conclusions & Challenges

- Interconnects are a critical enabling abstraction in MPSoC
- Interconnects play a very large and increasing role in delay, energy, and design effort.
- Interconnect can be solved simultaneously at the micro-architectural, circuit and process levels
- Aggressive circuit and signaling techniques show promise with minimal architectural impact
- CAD support needed, especially
  - early estimation for architecture and floorplanning
  - final verification in the presence of uncertainties



## *VLSI Interconnects: A Design Perspective,*

W. Burleson and A. Maheshwari

Morgan-Kaufmann. 2005(6)

- 400-page textbook with HW problems, covering:
  - History (both off-chip and on-chip)
  - Process (metallization, dielectrics, etc.)
  - Architecture (processor, ASIC, FPGA, memory)
  - Theoretical models (graph, information-theoretic)
  - Wire models (R,C,L,M,...)
  - Circuits (repeaters, boosters, sense-amps, etc.)
  - CAD (estimation, synthesis, optimization)
  - Case Studies (buses, memories, ASIC, FPGA)
  - Future (nano, optical, wireless, etc.)

# UMASS Interconnect Circuit Design Group

- Students:

- Vishak Venkatraman (internship at CRL, PhD 06)
- Jinwook Jang (MS 05, PhD 08)
- Sheng Xu (new Sept 04)
- Ibis Benito (new Jan 05)
- Atul Maheshwari (now at Intel)
- Matt Heath (now at Intel)
- Aiyappan Natarajan (now at AMD)
- Vijay Shankar (now at Qualcomm)
- Anki Nalamalpu (now at Intel)

- Collaborators:

- Sandip Kundu (UMASS/Intel/IBM)
- Russ Tessier, Israel Koren, Aura Ganz (UMASS)
- Shubu Mukherjee, Rich Watson, (Intel MMDC)
- SRC liasons (Intel CRL, Freescale)

- Alums

- Manoj Sinha (now at Micron)
- Chris Cowell (now at Intel)
- Sriram Srinivasan (now at AMD)
- Andrew Laffely (now Prof at USAFA)
- Srividya Srinivasaragavan (now at Intel)